Argonne NATIONAL LABORATORY

Magellan Cloud at ALCF

Susan Coghlan

Argonne Leadership Computing Facility

@

U.S. DEPARTMENT OF ENERGY

Magellan Background

- Joint project between Argonne Leadership Computing Facility (ALCF) and National Energy Research Scientific Computing Center (NERSC)
- Funded by DOE under the American Recovery and Reinvestment Act (ARRA)
- Why did ASCR fund Magellan?
  - ASCR workshop (http://www.sc.doe.gov/ascr/ProgramDocuments/Docs/MidrangeReportFinal.pdf) revealed need for mid-range computing
  - Although it is not part of ASCR's mission, midrange computing, and the associated data management play a vital and growing role in advancing science disciplines where capacity is as important as capabilty – Dan Hitchcock, ARRA project Briefing to ASCAC, Aug 11, 2009

Magellan Goal: To Explore!

- Discover science applications and user communities well suited for cloud computing
- Understand the deployment and support issues required to build large science clouds
  - Is it cost effective and practical to operate science clouds?
  - How can commercial clouds be leveraged? (stretch)
- Pioneer new cloud software and infrastructure that can better meet the needs of science
- Investigate how software as a service from the clouds can support science
- Examine how cloud computing can support data-intensive science
- Explore the challenges for security in a virtualized cloud environment

Unique Characteristics of Magellan

- High Speed, Low Latency Interconnect
  - QDR Infiniband Connection to All Nodes
- High Performance Storage
  - Solid State Storage
  - High Performance Parallel Filesystem
- High Bandwidth Wide Area Networking
  - Direct Connection to 100Gbps ANI
- Tuned Middleware and Scientific Software

Magellan at NERSC
Purpose-built for Science Applications

720 nodes, 5760 cores in 9 Scalable Units (SUs) ➔ 61.9 Teraflops
SU = IBM iDataplex rack with 640 Intel Nehalem cores

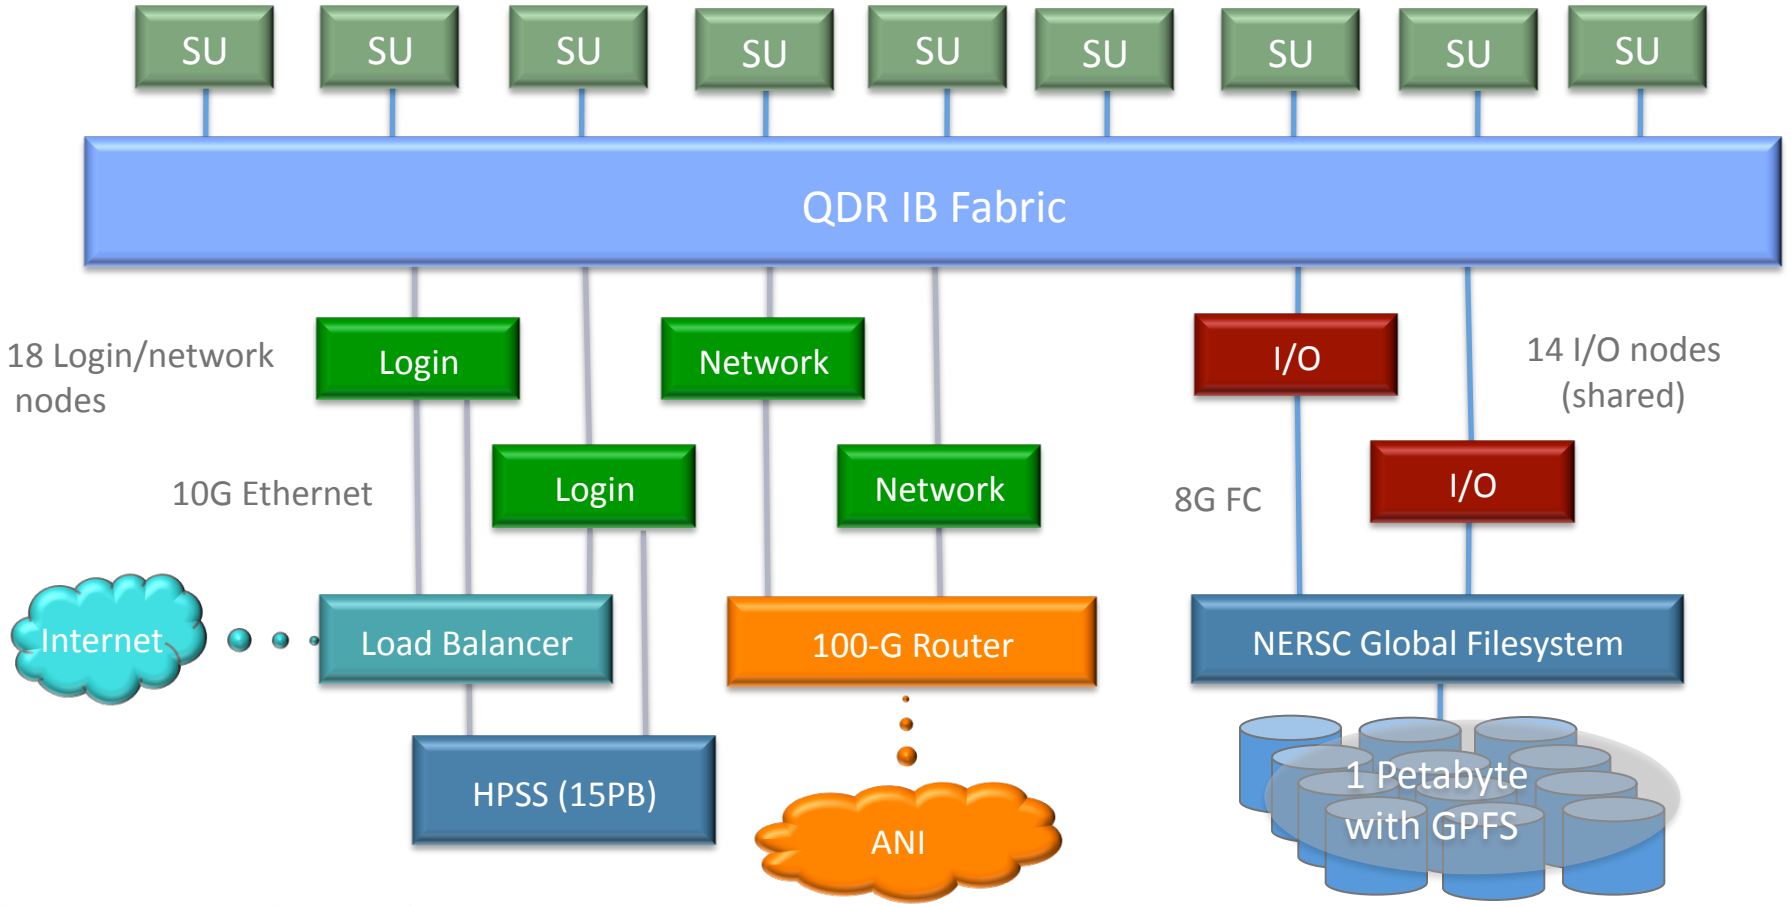

Flexible and dynamic scheduling of resources

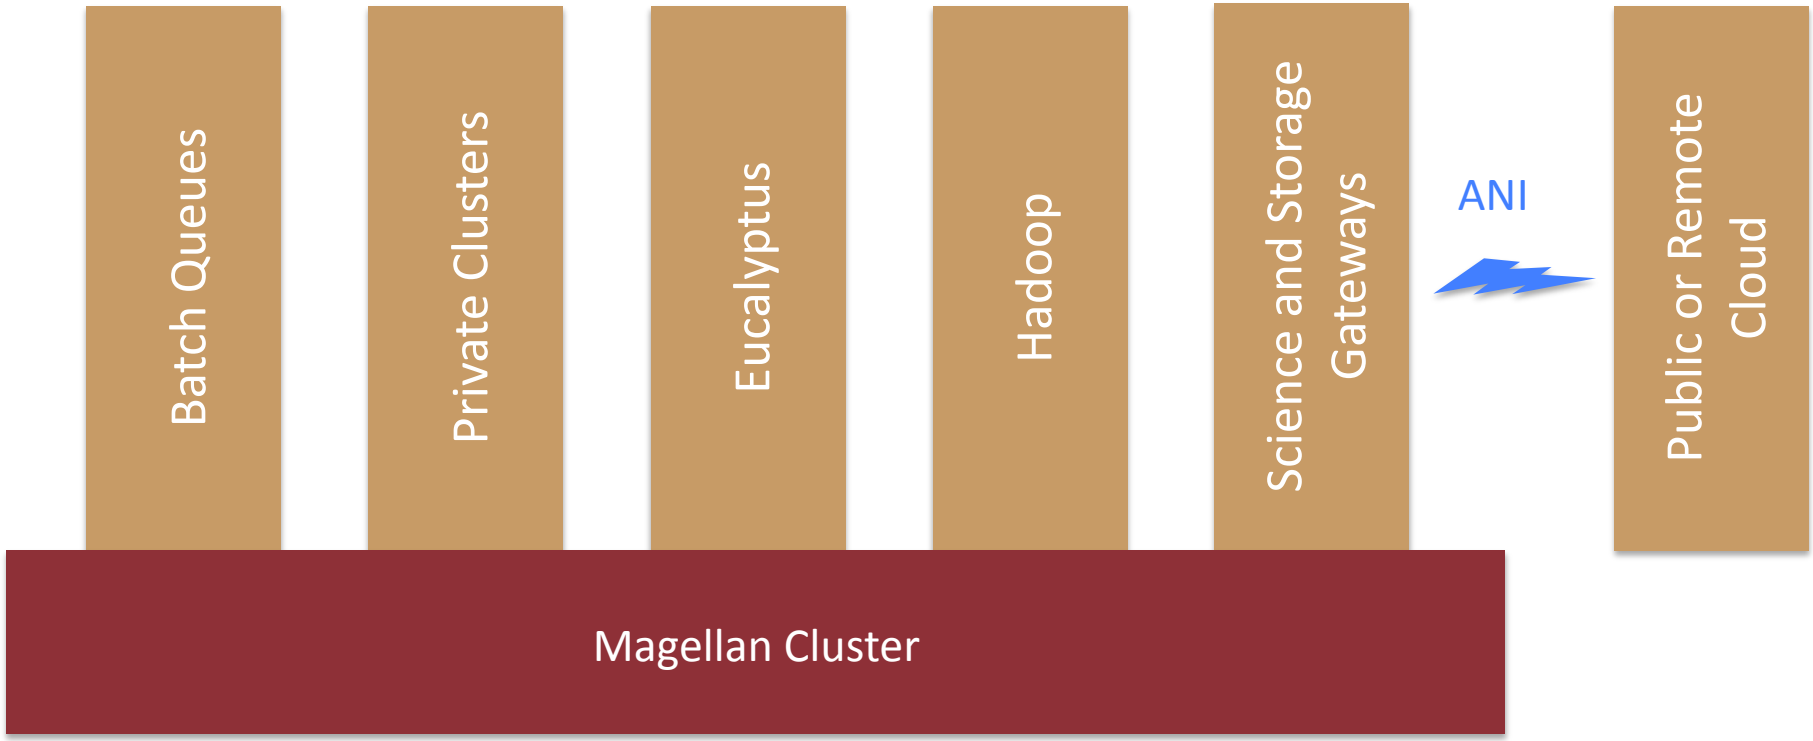

- Runtime provisioning of software images
- Rolling upgrades can improve availability
- Choice of local or remote cloud

Argonne Magellan Cloud Hardware – Phase 1 Spring 2010

Compute

504 Compute Nodes
- Nehalem Dual quad-core 2.66GHz
- 24GB RAM, 500GB Disk
- QDR IB link

Totals
- 4032 Cores, 40TF Peak
- 12TB RAM, 250TB Disk

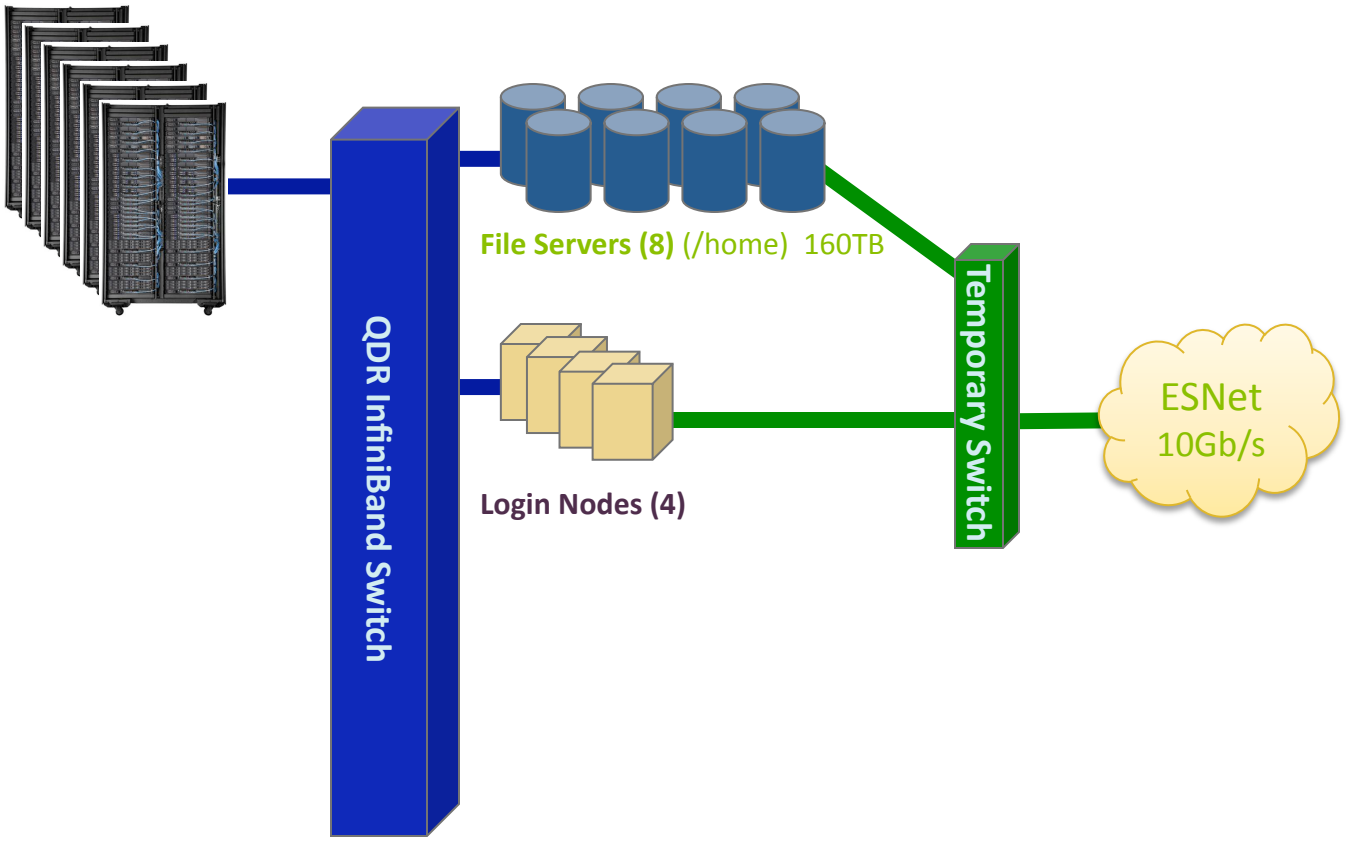

Argonne Magellan Cloud Hardware – Phase 2
Late Summer 2010

Compute

504 Compute Nodes
- Nehalem Dual quad-core 2.66GHz
- 24GB RAM, 500GB Disk
- QDR IB link

Totals
- 4032 Cores, 40TF Peak
- 12TB RAM, 250TB Disk

Active Storage

~100 Compute/Storage Nodes
~10TB FLASH/SSD Storage
~500TB Disk Storage

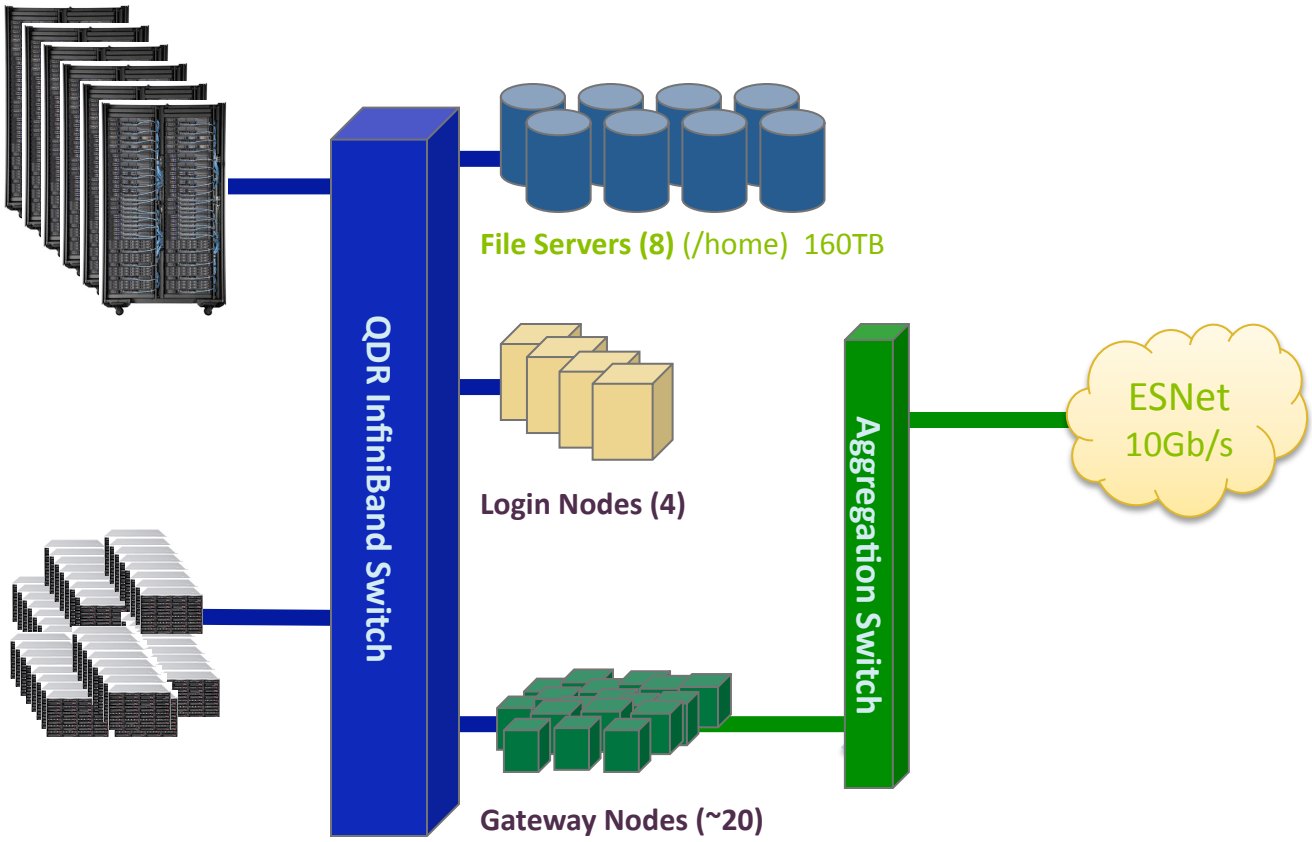

Argonne Magellan Cloud Hardware – Final January 2011

Compute

504 Compute Nodes
- Nehalem Dual quad-core 2.66GHz
- 24GB RAM, 500GB Disk
- QDR IB link

Totals
- 4032 Cores, 40TF Peak
- 12TB RAM, 250TB Disk

Active Storage

~100 Compute/Storage Nodes
~10TB FLASH/SSD Storage
~500TB Disk Storage

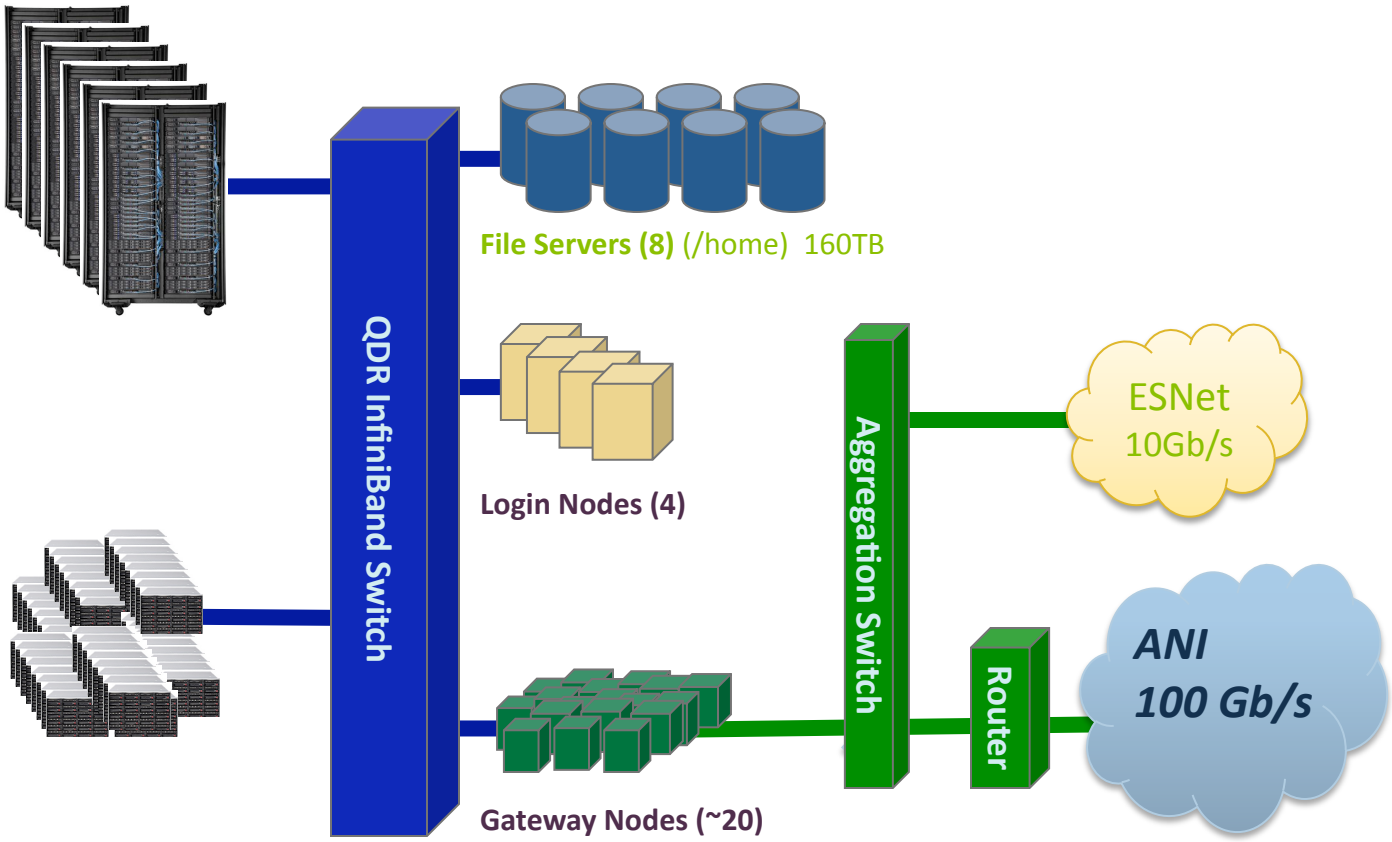

Argonne Magellan Software Architecture

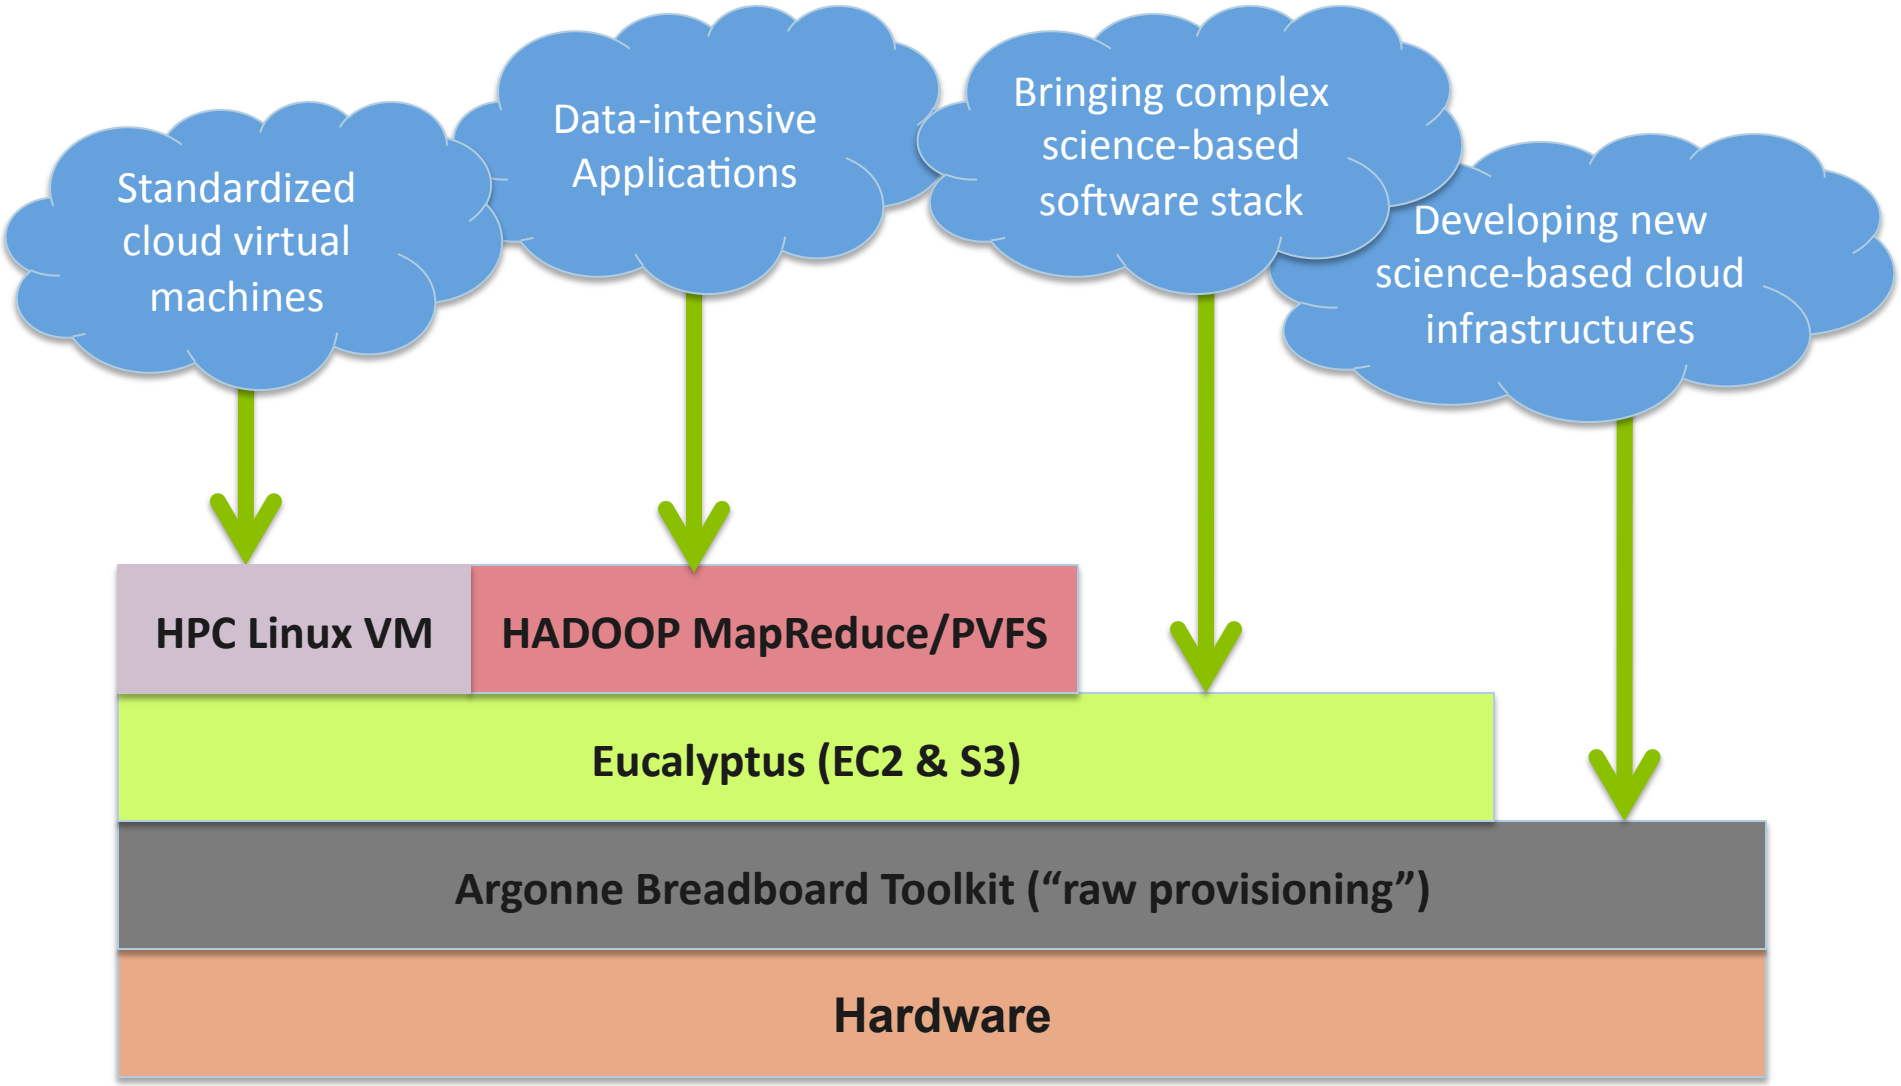

Magellan Users

- Science teams
  - Explore using the cloud to accomplish science
- ANI 100Gb/s project teams
  - Add unique capabilities to cloud computing with fast networking
- Application Developers
  - Develop applications and ensemble pipelines well-suited for clouds
- System software developers
  - Develop software that provides unique capabilities for science applications running in the cloud
- System manager
  - Integrate transient resources into local production offering
- System operations
  - Provide secure environment and coordinate ALCF/NERSC services

ALCF Magellan High-Level Time Line

- January 2010 – Core System Delivered
- March 2010 – Open to Early Users, Welcome to Magellan Day (3/23)
- April 2010 – Open to regular users
- September 2010 – Phase 2 hardware delivered
- November 2010 – Phase 2 hardware open to regular users
- January 2011 – 100Gbit deployed

We're Looking for a Few Good Users

- Projects That
  - Exercise Unique Properties of Magellan and Clouds
  - Leverage Research at ALCF and MCS
  - Are Not Well-Served by Current Resources
- People That
  - Tolerate Downtime and Configuration Changes
  - Provide Feedback and Cooperative Experimentation
  - Have the Time to be Part of the Magellan Team

How to Get Involved

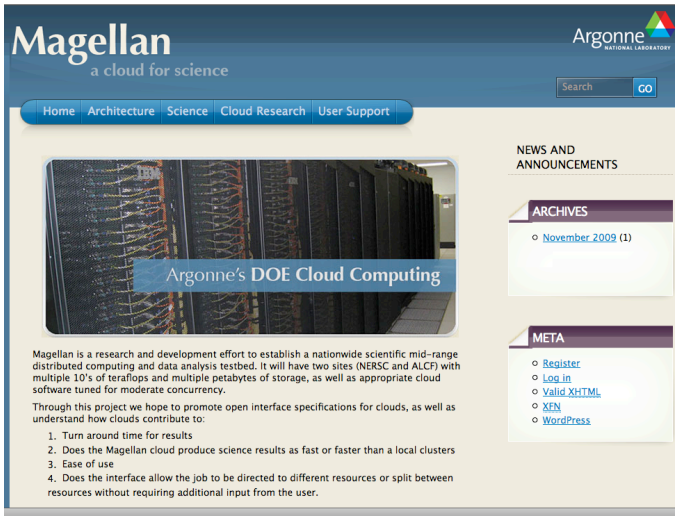

- http://magellan.alcf.anl.gov
- Apply for a Project
  - Application Description
  - Software Description
  - Interest in Specific Cloud Capabilities
  - Plans for Collaboration with ANL Cloud Research
  - Allocation Request
- Discussion List: magellan@alcf.anl.gov
- Support: magellan-support@alcf.anl.gov